Compiled by Dr T Oberholster
AI-900 version: October 2021

Website & Blog
aRtVerse Shop
Amazon Books
Teepublic Shirts
SpreadShirt Apparel

What is ML?

- Process of using mathematical models of data to help a computer learn without direct instruction
- Subset of Artificial Intelligence (AI)

How does it work?

- Algorithms identify patterns within data, which are used to create a data model that can make predictions with new (related) data.
- Type of predictive analytics with real-time updates
- Deep learning: specialized ML train models to mimic human reasoning (such as a neural network, which is modelled after the human brain)

Benefits

- Uncover insights: ML identifies patterns (data story) in structured & unstructured data
- Enhance user experience: Adaptive interfaces, targeted content, voice-enabled virtual assistants help optimise customer experience
- Anticipate customer behaviour: Mining of customer-related data to identify patterns/behaviour and optimize product recomendations/experience
- Improve data integrity: ML improves over time with data mining
- Reduce risk: Fraud tactics change, ML keeps pace to monitor & ID new patterns & catch attempts
- Lower costs: Process automation frees up time and resources for human teams

ML as a Service

- Data scientists & developers build, train and deploy machine learning models
- Accelerated time to market with MLOps (DevOps for ML)
- All skill levels
- Responsible ML innovation (protect data)
- Cost management with resources

Machine Learning

Pre-Built Services

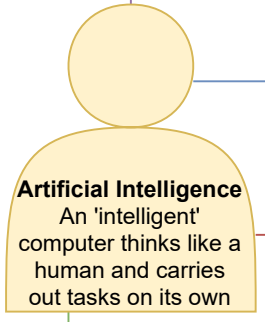

Custom Services

MS ML Services

Anomaly Detection

- Auto detect errors or unusual activity in a system

Computer Vision

- Software interpret world visually through cameras, video and images

Natural Language Processing

- Computer interpret written and spoken languages & responds in kind

Conversational AI

- Software 'agent' participate in a conversation

Azure Machine Learning Service

- Cloud based platform to create, manage and publish ML models

Create Machine Learning Models

- AutoML
  - Non-experts can quickly create effective ML models
- ML Designer
  - Graphical interface enabling no-code development
- Data & compute management
  - Cloud based storage and compute resources to create code at scale
- Pipelines
  - Orchestrate model training, deployment and management

ML Technique to forsee problems

- Analyses data over time (time series data) & identifies unusual changes
- Application Programming Interface (API)
  - Univariate (single variable/metric)
  - Multivariate (groups of metrics)
- Auto identify and apply best-fit model regardless of industry, scenario or data volume
- Determines boundaries for anomaly detection, expected values and which data points are anomalies
- No prior ML knowledge required
- RESTful API integration with other services

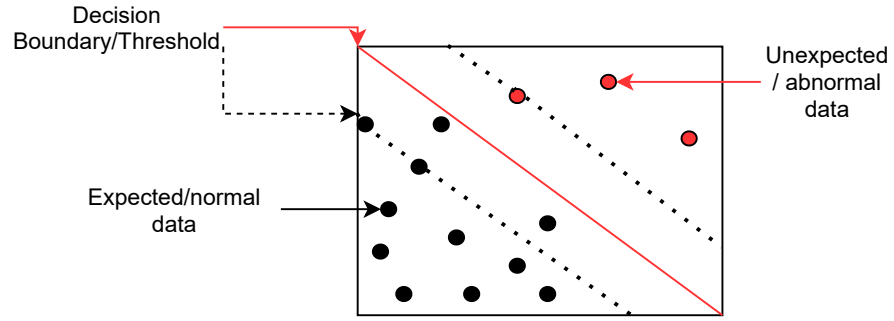

Concepts

Features

Azure Anomaly Detector

Use cases

Anomaly detection in real-time

- Steaming data

Batch detect anomalies

- Detect any anomalies in data

Batch detect change points

- Detect trend change points in data

Additional information

- Details about anomalies (expected values, boundaries and positions)

Adjust anomaly detection boundaries

- Increase or decrease sensitivity to anomalies to better fit the data

- Suspicious credit card transactions
- Track production line failures
- Proactive warning system for mechanical failures before they occur

aRt Verse
ART + SCIENCE + NATURE

Compiled by Dr T Oberholster
AI-900 version: October 2021

Website & Blog
aRtVerse Shop
Amazon Books
Teepublic Shirts
SpreadShirt Apparel

Compiled by Dr T Oberholster
AI-900 version: October 2021

aRtVerse
ART + SCIENCE + NATURE

Website & Blog
aRtVerse Shop
Amazon Books
Teepublic Shirts
SpreadShirt Apparel

ML Technique to extract information

- Models based on visual input from cameras, videos or images
- Subsets
  - Image classification: Single subject classified based on contents
  - Object detection: Multiple subjects
  - Semantic segmentation: Multi-subject pixel classification
  - Image analyses: Descriptive tags for content
  - Face detection: Recognize/analyse human faces
  - Optical character recognition (OCR): Detect and read text in images/documents
- To AI an image is an array of pixel values (i.e., numeric features to train models)
  - Needs lots of data
  - Is time consuming
  - MS AZ Computer Vision provides pre-built models

Concepts

Cognitive Services

Azure Computer Vision

Use cases

- "Seeing AI" describes objects, people and text for visually impaired people

- **Content organization:** people, objects organized into categories for photo storage/social media applications
- **Text extraction:** Analyse images and PDF documents that contain text & extract text into structued format
- **Spatial Analyses:** Identify people/objects in space and map their movement (people and vehicles, people in retail stores/queue wait times)
- **Describe an image**: Analyse image, evaluate object & generate human-readable descriptive phrase (caption). May return multiple results with confidence scores
- **Tagging visual features:** Recognize objects & suggest tags, metadata, summarise attributes for image indexing & search
- **Detect objects:** Similar to tagging, ID objects & return bounding box coordinates
- **Detect brands:** ID commercial global logos & return brand name with confidence score & bounding box coordinates
- **Detect faces**: Detect & analyse faces (age, bounding box coordinates for location in image). Face Service provides comprehensive facial analyses
- **Categorize images:** Based on contents, as a parent/child hierarchy with limited categories
- **Domain-specific content:** ID celebrities (sport, entertainment, business) and Landmarks (global locations/buildings)
- **Optical Character Recognition:** Detect printed/handwritten text in images
- **Detect image types:** Clipart, line drawing
- **Image colour schemes:** Foreground, background/overall
- **Generate thumbnails:** Small versions of images
- **Moderate content:** Detect adult content/violent scenes

Services

MS AZ Computer Vision

- Analyse images/videos
- Extract/Create tags, objects, text & descriptions
- No ML expertise required

MS AZ Custom Vision

- Train custom image classification and/or object detection models with own images

MS AZ Face

- Build facial detection and recognition solutions

MS AZ Form Recognizer

- Extract information from scanned forms and invoices

Resourcing for Service

Provision a Resource

- Computer Vision (CV) resource: Use CV only, or manage access/utilization/tracking/billing of resource separately
- Cognitive Serices (CS) resource: Combine services, or manage access/utilization/tracking/billing of resource together
- Provice client application authenticaion key and HTTP address endpoint for ML service consumption

ML Technique to ID real-world objects

- Models predict category/class of object
  - Set of inputs (features) to predict labels with probability scores
  - Trained on pixel values (feautres) and labels of categorized images
- Can be combined with deep learning techniques and convolutional neural networks
  - Complicated and requires expertise to build from scratch
  - Available in Azure Custom Vision Cognitive Services

MS AZ Computer Vision

- Subset of Analyze Image of Computer Vision
- No need to manage individual servers for ML modelling and classification
- Auto-scale
- Azure managed identities
- Regional resiliency

- Businesses that need to process and classify images
- Product identification
- **Disaster investigation:** Infrastructure management
- **Medical diagnosis:** Xrays or MRI
- **Anomaly detection for defects**: Manufacturing
- **Image classification:** Fashion website, screenshots, insurance claims

Step 4: Publish model to client application for consumption

- Provide the following details for deployment
  - Project ID: Unique ID of Custom Vision project
  - Model name: Name of model
  - Prediction endpoint: HTTP address of endpoints for prediction resource and published model
  - Prediction key: Authentication key for published model

aRt Verse
ART + SCIENCE + NATURE

Compiled by Dr T Oberholster
AI-900 version: October 2021

Website & Blog
aRtVerse Shop
Amazon Books
Teepublic Shirts
SpreadShirt Apparel

Concepts

Services

Use Cases

Build an ML Model

Cognitive Services: Computer Vision

Azure Image Classification

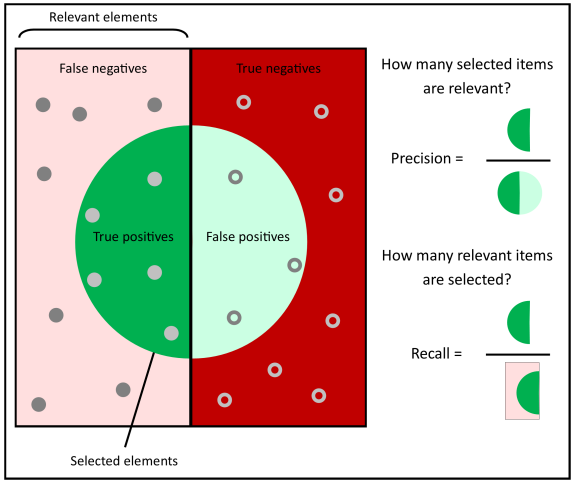

Step 1: Provision a Resource

- Computer Vision (CV) resource: Dedicated training/ prediction or both resources. Use CV only, or manage access/ utilization/ tracking/ billing of resource separately. Cannot use authentication key/ service endpoint interchangeably.
- Cognitive Serices (CS) resource: Dedicated training/ prediction or both resources. Combine Custom Vision or other services, or manage access/ utilization/ tracking/ billing of resource together. One endpoint and authentication key for resources.
- Provide client application authenticaion key and HTTP address endpoint for ML service consumption

Step 2: Train the Model with labelled historical data

- Iteratively train the model in the Azure portal or with code via the software development kits (SDKs)
- Hold out some data for model evaluation (perfomance testing)

Step 3: Evaluate model performance with metrics

- Determine model perfomance by comparing predictions with actual (hold out) data.
- Precision
  - Percentage of identified classification that were correct. Model identified 100 images belonging to category A, and 90 were indeed A then precision = 90%.
- Recall
  - Percentage of actual classifications that were correct. Model identified 80 images belonging to category A, but there are 100 images of A then recall = 80%
- Average Precision (AP)
  - Overall metric that includes precision and recall

Compiled by Dr T Oberholster
AI-900 version: October 2021

aRtVerse
ART + SCIENCE + NATURE

Website & Blog
aRtVerse Shop
Amazon Books
Teepublic Shirts
SpreadShirt Apparel

Azure Object Detection

Cognitive Services: Computer Vision

ML Technique to detect common objects in images

- Model recognizes individual objects in an image
- Class of each object in image
- Probability score of classification (confidence of prediction)
- Coordinate bounding box for each object
- Similar to tagging
- Determine multiple instances of same object

Concepts

Services

MS AZ Computer Vision

- Subset of Analyze Image of Computer Vision
- Detect API
  - Applies tags in images
- Tag API
  - Include contextual terms (i.e., indoors)

Use Cases

- Safety evaluation of fire equipment
- Lane assist capabilities in self-driving cars
- Medical imaging (MRI or X-ray)

Limitations

- Small objects (5% of image)
- Arranged closely together (stacked objects)
- No brand/product name IDs

Build an ML Model

Step 1: Provision a Resource

- Computer Vision (CV) resource: Dedicated training/ prediction or both resources. Use CV only, or manage access/ utilization/ tracking/ billing of resource separately. Cannot use authentication key/ service endpoint interchangeably.
- Cognitive Services (CS) resource: Dedicated training/ prediction or both resources. Combine Custom Vision or other services, or manage access/ utilization/ tracking/ billing of resource together. One endpoint and authentication key for resources.
- Provide client application authenticaion key and HTTP address endpoint for ML service consumption

Step 2: Image tagging

- Tag classes and bounding box coordinates
- Time consuming
- Custom Vision GUI:
  - Suggests areas of discrete objects for tagging, drag area sizes.
  - Smart tagging suggests classes for additional training images.

Step 3: Model training

- Iteratively train model in portal or with code via software development kits (SDKs)
- Hold out some data for model evaluation

Step 4: Evaluate model performance with metrics

- Determine model perfomance by comparing predictions with actual (hold out) data.
- Precision
  - Percentage of class predictions that were correct. Model identified 100 images belonging to category A, and 90 were indeed A then precision = 90%.
- Recall
  - Percentage of class prediction correctly identified. Model identified 80 images belonging to category A, but there are 100 images of A then recall = 80%
- Average Precision (AP)
  - Overall metric that includes precision and recall

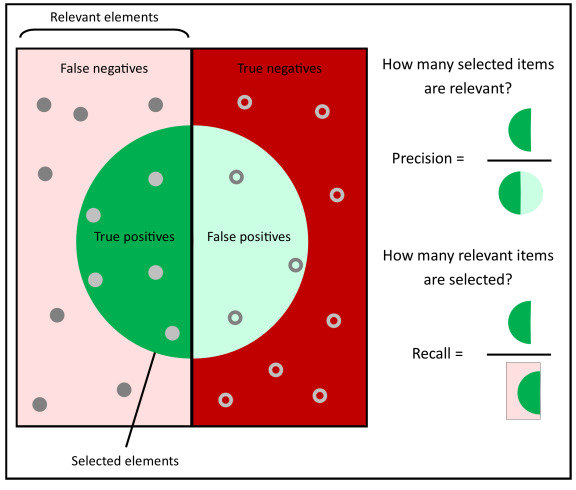

Step 5: Publish model to client application for consumption

- Provide the following details for deployment
  - Project ID: Unique ID of Custom Vision project
  - Model name: Name of model
  - Prediction endpoint: HTTP address of endpoints for prediction resource and published model
  - Prediction key: Authentication key for published model

Compiled by Dr T Oberholster
AI-900 version: October 2021

Website & Blog
aRtVerse Shop
Amazon Books
Teepublic Shirts
SpreadShirt Apparel

Azure Face Detection

Cognitive Services: Computer Vision

Concepts

ML Technique to analyze human faces

- Locate and analyse human faces in images/video
- **Face detection:** ID regions of image with human face & return bounding box coordinates
- **Facial analysis:** Return additional information (facial landmarks: nose, eyes, eyebrows and lips) and used to infer age or emotional state, wearing glasses/masks
- **Facial recognition:** Identify known individuals from facial features. Use multiple images to train model for identification

Services

MS AZ Computer Vision

- Subset of Analyze Image of Computer Vision
- No ML expertise required
- Computer Vision: Basic analyses (age)
- Video Indexer: Faces in videos
- Face: Widest range of capailities
  - Identify API: One-to-many face match
  - Verify API: One-to-one face match
  - Group API: Similartiy face match

Use Cases

- **Security:** unlock doors or smart devices (toucheless access control)
- **Social media:** tag friends in photographs
- **Intelligent monitoring:** drivers looking at road/mobile device/tired
- **Advertising:** direct advertisements to appropriate demographic audience
- **Missing persons:** public camera systems, facial recognition to ID missing person in the frame
- **Identity validation/verification:** Port of entry determine if person holds correct permit
- **Similar faces:** Find similar faces during a search

Build an ML Model

Step 1: Provision a Resource

- Face (CV) resource: Dedicated training/ prediction or both resources. Use CV only, or manage access/ utilization/ tracking/ billing of resource separately. Cannot use authentication key/ service endpoint interchangeably.
- Cognitive Services (CS) resource: Dedicated training/ prediction or both resources. Combine Custom Vision or other services, or manage access/ utilization/ tracking/ billing of resource together. One endpoint and authentication key for resources.
- Provide client application authenticaion key and HTTP address endpoint for ML service consumption

Features

Face Supports

- Face detection, verification, find similar faces, group faces based on similarities & identify people
- Age, blur, emotion, light exposure, facial hair, glasses, hair, head pose, makeup, visual noise/clarity, occlusion, smile

Performance

Improve accuracy

- JPEG, PNG, GIF and BMP images
- < 6 MB
- 36 x 36 to 4096 x 4096 pixel faces
- Limit extreme face angles, occlusions (objects blocking face), preferably full-frontal face image

Compiled by Dr T Oberholster
AI-900 version: October 2021

aRtVerse
ART + SCIENCE + NATURE

Website & Blog
aRtVerse Shop
Amazon Books
Teepublic Shirts
SpreadShirt Apparel

ML Technique interpret written/printed text

- Intersection of Computer Vision and Natural Language Processing
- Optical character recognition (OCR)
  - Models recognize individual shapes as letters, numbers and punctuation
- Machine Learning Comprehension
  - AI system reads and understands semantic context

Concepts

MS AZ Computer Vision

- Subset of Analyze Image of Computer Vision
- OCR API
- Read API
- Form Recognizer API

Services

- Note taking
- Digitizing forms
- Scanning bank checks
- Expense claim
- Sales receipts

Use Cases

Cognitive Services: Computer Vision

Optical Character Recognition

Build an ML Model

Step 1: Provision a Resource

- Computer Vision (CV) resource: Dedicated training/ prediction or both resources. Use CV only, or manage access/ utilization/ tracking/ billing of resource separately. Cannot use authentication key/ service endpoint interchangeably.
- Cognitive Services (CS) resource: Dedicated training/ prediction or both resources. Combine Custom Vision or other services (Text Analytics, Translator Text), or manage access/ utilization/ tracking/ billing of resource together. One endpoint and authentication key for resources.
- Provide client application authenticaion key and HTTP address endpoint for ML service consumption

OCR API

- Quick extraction of small amounts of text in images
  - Synchronous and multi-language
- Information hierarchy
  - Regions that contain text
  - Lines of text in region
  - Words of each line of text
- Returns bounding box coordinates of region, line or word
- OCR generates false positives with text-dominated images

Text Limited

Text Dominant

Read API

- Optimized for text-dominant or images with visual noise
- Scanned documents
- Auto determines best model to use
- Asynchronous
- Submit images, receive operation ID response
  - Use operation ID to check status
  - Receive results of operation
    - Pages (one per page of text, with size & orientation)
    - Lines
    - Words

OCR for receipts and invoices

- Process electronically for expense claims or digitization
  - Name, address and telephone number of merchant
  - Date and time of purchase
  - Quantity and price of item
  - Subtotal, tax and total amounts
- Custom models
  - Extract key/value pairs and table data from forms
  - Trained with own data to tailor models for specific forms
  - Five samples to start

Best performance

- JPEG, PNG, BMP, PDF or TIFF formats
- File size < 50 MB
- 50 x 50 to 10000 x 10000 px
- PDFs no larger than 17 x 17 inches

Automated document processing

Form Recognizer

- Automate data processing from documents (forms, invoices and receipts) extract text, key-value pairs, tables, and structures from documents
- Combines ORC with predictive models to
  - Match field names to values
  - Process tables of data
  - Identify types of fields (dates, addresses, totals)
- Pre-built receipt model for sales receipts
  - Common receipts (USA) in English
  - Restaurants, retail locations and gas stations
  - Time and date, merchant, taxes, totals, other pertinent info
  - All text recognized is returned
- Free tier: First two pages processed (PDF or TIFF)

Compiled by Dr T Oberholster
AI-900 version: October 2021

aRtVerse
ART + SCIENCE + NATURE

Website & Blog
aRtVerse Shop
Amazon Books
Teepublic Shirts
SpreadShirt Apparel

ML Technique mines insights from unstructured text

- Understands written and spoken language
- Analyse & interpret text with AI
- Interpret spoken language & synthesise speech response
- Translate spoken or written language
- Interpret commands and take actions
- Evaluate different aspects of document/phrase to gain insights into the context of the text.

Concepts

Build an ML Model

Step 1: Provision a Resource

- Text Analytics/Translator Text/Speech/LUIS resource: Dedicated training/ prediction or both resources. Use this resource only, or manage access/ utilization/ tracking/ billing of resource separately. Cannot use authentication key/ service endpoint interchangeably.
- Cognitive Services (CS) resource: Dedicated training/ prediction or both resources. Combine Custom Vision or other services (Text Analytics, Translator Text), or manage access/ utilization/ tracking/ billing of resource together. One endpoint and authentication key for resources.
- Provide client application authenticaion key and HTTP address endpoint for ML service consumption

MS AZ Computer Vision

- Subset of Analyze Image of Computer Vision
- Pretrained NLP ML models and services
  - **Text Analytics:** Analyse text documents and extract phrases, detect entities (places, dates, people), and evaluate sentiment (positive/negative)
  - **Translator Text:** Translate +60 languages
  - **Speech:** Recognize, synthesise, and translate speech
  - **Language Understanding Intelligent Service (LUIS):** Train language model to understand spoken or text-based commands

Services

Cognitive Services for Language

Speech

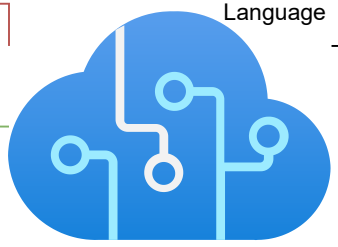

Natural Language Processing & Conversational AI

Use Cases

- ID and categorize concepts
- ID main points in unstructured text
- Understand customer perception
- Process unstructed medical/clincal data
- Conversational interface with FAQs/manuals/website
- Automate workflow by classifying unstrucuted text

Text

Speech service

- Speech-to-text
- Text-to-speech
- Speech translation
- Speaker recognition

- AI accepts vocal commands and returns spoken responses (home and auto systems)
  - **Speech recognition:** Ability to detect & interpret spoken language
  - **Speech synthesis:** Ability to generate spoken output

Text Analytics Service

- Statistical analysis of terms (remove stop words) with frequency analyses (counting words) for clues on main topic
- Extend frequency analyses to multi-term phrases
  - N-grams, a two-word phrase = bigram and three-word phrase = trigram
- Stemming or lemmatization algorithms to normalize words before counting (only use root words)
- Analyse sentences with linguistic structure rules (noun, verbs, and adjectives)
- Encode words/terms as numeric features to train ML model.
  - Classify document based on terms inside.
  - Sentiment analyses to classify document as positive/negative
  - Create vectorized models to capture semantic relationships between words in n-dimensional space. Related terms are grouped closer together.

Speech recognition (Speech service)

- Transcribes spoken word (recorded or live) into data for processing
- Speech patterns analysed & mapped to words using multiple models
  - **Acoustic model:** Converts audio signal into phonemes (represent specific sounds)
  - **Language model:** Maps phonemes to words (using statistical model and confidence scores)
- Words converted to text
  - Closed captions for videos
  - Transcript of phone call or meeting
  - Automated note dictation
  - Determine intent of user for processing

Text Analytics Service

- **Language detection:**
  - Language name
  - ISO 6391 language code
  - Confidence score of detection
  - Focus on predominant language in the text in mixed language submissions
  - Ambiguous language (keyboard emoji) evaluated as unknown with NaN score
- **Sentiment analysis:**
  - Evaluate & return sentiment scores & labels for sentences
  - Range 0 – 1, 1 = positive, 0.5 = neural/indeterminate, 0 = negative
  - Indeterminate may be due to insufficient context, no sentence structure or incorrect language used (= 0.5)
- **Key phrase extraction:**
  - Evaluate text & identify main topic(s)
  - Entity recognition
  - List of entities (people, places, dates, objects)
    - Categories (person, location, organization, quantity, datetime, URL, email, phone number, IP address)
    - Subtypes (number, percentage, ordinal, age, currency, dimension, temperature, date, time, daterange, timerange, duration or set)
  - Entity linking to disambiguate entities using a specific reference.

Speech synthesis (Speech service)

- 'Reverse' of speech recognition
- Vocalising data, text converted to speech
  - Text to be spoken
  - Voice used to vocalize speech
- Systems tokenizes text (broken down into word) & assigns phonetic sounds to words
  - Phonetic transcription broken down into prosodic units (phrases, clauses, or sentences) to create phonemes to convert to audio and synthesised to words using a voice
  - Voice determines pitch and timbre in audio wave form to output to speaker or written to file
- Used for
  - Generating spoken responses to user input
  - Voice menus for telephone systems
  - Read email or text messages (hands free)
  - Broadcast announcements (public locations, stations, or airports)

Speech-to-Text API (Speech service)

- Convert text to audible speech, played through computer speaker or saved to file
- Pre-built voices with multi-language & regional pronunciation support
- Standard and neural voices (leverage neural networks to create intonation and more natural sounding voices)
- Develop custom voices

Translate Text and Speech

- Automated or machine translation allows people from different culture and/or geographical location to collaborate
- Multi-language support for text translation, transliteration, language detection, and dictionaries.
- **Literal translation:** Each word translated to corresponding word. Issues include no equivalent words or changes in phrase/meaning/context.
- **Semantic translation:** AI understands context of the words used and returns a more accurate translation where grammar, formal/informal and colloquialisms are considered.
- **Text translation:** Translate documents from one language to another (email, web pages)
- **Speech translation:** Translate between spoken languages, directly (speech-to-speech) and via text (speech-to-text-translation).

Translator Text Service and Speech Service

- For applications, websites, or tools
  - Neural Machine Translation (NMT) analyses semantic context of text and renders more accurate/complete translation
  - 60+ languages support
  - Provide ISO 639-1 language or 3166-1 cultural codes for languages being translated
  - One-to-one and one-to-many translations

- Optional configurations
  - **Profanity filtering:** Culture specific profanity removed by marking or omitting results
  - **Selective translation:** Tag content not to be translated (code, brand, name, word/phrase)
- **Speech service:** speech-to-text and speech-to-speech translation
  - Speech-to-text, Text-to-speech, and Speech translation
  - **Speech translation:** translate spoken word from streaming audio (live or recorded) for real-time closed captioning or two-way communication
    - One-to-one or One-to-many translations
    - Specify language and/or culture code format to consider localized pronunciation and linguistic idioms

Language Understanding

- MS Azure Language Understanding Intelligent Service (LUIS)
  - **Utterances:** Question/phrase/command
    - Different ways of saying the same thing
    - Used to train the model
  - **Entities:** An object/person/place to which is being referred
    - To which an action is applied by the model
    - List: Entities in hierarchical list/sublist and includes synonyms
    - RegEx: Designed as regular expression (based on a pattern of alphanumerical characters)
    - Pattern.any: Patterns for complex entities that may be hard to extract from sample utterances
  - **Intents:** Purpose/goal as determined by the utterance
    - Action for model to execute on an entity
    - There are numerous utterances, grouped by intent and related to entities
    - None intent is a fallback for unknown utterances with generic responses
- After model training (authoring), the model predicts the intent and entity based on the utterance to respond appropriately (prediction).
- Two resources made for Authoring and Prediction to separate tracking of resource utilization
  - Authoring
    - Author and train LU application with utterances, intents & entities
    - Prebuilt domains with predefined intents and entities for common scenarios
    - Custom intents and entities and map to utterances
    - Create model through web-based LU portal or with code
    - After training responses can be reviewed and updated.
  - Predicting
    - Trained LU model is published to application for consumption
    - Connect to a client application endpoint, specify authentication key and submit user input for predictions and return an action based on intent

Conversational AI

- Bots manage and participate in conversations with humans
- Dialogs through website, email, social media, messaging, phone calls and other channels
- Customer product/services, reservation systems, health care consultants, home automation

QnA Maker Service

- Build a knowledge base of question-and-answer pairs, usually with built-in NLP for multiple phrased examples but same semantic meaning.
- Code with REST API or SDK or use QnA Marker portal
- QnA pairs created from existing FAQ or page, imported from chit-chat data source, or manually created or combo
    - Alternative phrasing consolidates questions with same meaning
    - Train knowledge base model with literal QnA then test with built-in interface and review answers
- Publish to client applications with knowledge base ID, endpoint, and authorization key
    - Deliver knowledge base via a bot, created in MS Bot Framework SDK to write conversational flow and integration code
    - Or use automatic bot creation in QnA Maker via Azure Bot Service
- Manage the bot in Azure portal
    - Extend functionality with code (simple edit in Azure portal, comprehensive ~ download and edit locally and republish to Azure)
    - Test the bot's interactive interface
    - Config logging, analytics, and integration with other services
    - Connect bot to multiple channels, users interact through chat, email, Teams, and other media.

Azure Bot Service

- Develop, publish, and manage bots.
- Bot Framework to create bot, manage with Bot Service and integrate QnA/LUIS for interactions.

Challenges & risks

- **Bias:** Race or gender bias in model due to bias in data
- **Errors:** When system fails, drastic consequences
- **Breach:** Sensitive data leaked/stored insecurely
- **Suboptimal solutions:** Caters to specific user, not universal
- **Trust:** Users do not know about underlying data/assumptions
- **Liability:** Who is responsible when predictions are wrong

MS AI software development guided by 6 principles

- Ensures the provisioning of solutions and minimise unintended negative consequences
  - Fairness
  - Reliability & safety
  - Privacy & Security
  - Inclusiveness
  - Transparency
  - Accountability

Understand, protect, and control your data, models, and processes to build trusted solutions

- State-of-the-art technology enabling responsible ML development, deployment, and use.
- Responsible AI allows ML to be understandable & cultivates trust instead of it being a 'black box'

Principles

Concepts

In Depth

Azure Responsible AI Principles

aRtVerse
ART + SCIENCE + NATURE

Compiled by Dr T Oberholster
AI-900 version: October 2021

Website & Blog
aRtVerse Shop
Amazon Books
Teepublic Shirts
SpreadShirt Apparel

- **Fairness:** Treat all people fairly without bias/discrimination with regards to opportunities, resources, and information. Mitigated through interpretability and explainability of the model (Sociotechnical challenge)
- **Reliability & safety:** Does not cause harm, preferably even during failure or quantifies potential risk/harm. Mitigated through rigorous testing and development management.
  - Track and document ML process (datasets, models, code ect)
- **Privacy & Security:** Secure sensitive information and respect privacy. Mitigated through running model on user devices or using anomaly detection on systems.
  - Encrypt data, create noise to keep data confidential
- **Inclusiveness:** Empower and engage everyone. Mitigated through including diverse perspectives during model development.
- **Transparency:** Understandable; why was it created, how it works and what its limitations are. Mitigated through interpretability of the models.
  - Undestand model behaviour and features that impact predictions
- **Accountability:** Who is accountable for the system and its results. AI created in a framework of clear governance, ethical and legal standards.

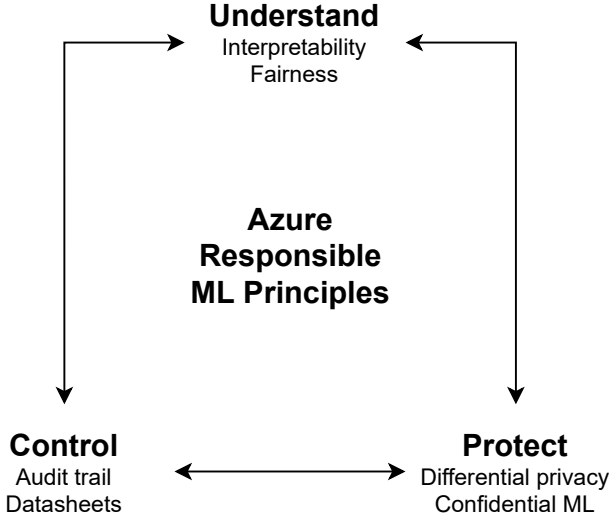

ML to detect patterns in historical data to make new predictions

- Train predictive models using historical data, statistics and mathematics to make predictions about unknown values.
- Function *f* operates on features x to calculate the predicted label (y)
  - *f*(x) = y (function fits features to label, function calculates y from x)
- Data requires pre-processing and specific operations to use - this depends on the model and training algorithm
  - Data Ingestion (gather data)
  - Data pre-processing (feature selection/identify x, clean data, impute missing, feature engineering/derive new features, normalization to same scale),
  - Train model and fit to data,
  - Predict unknown values.

aRtVerse
ART • SCIENCE • NATURE

Compiled by Dr T Oberholster
AI-900 version: October 2021

Website & Blog
aRtVerse Shop
Amazon Books
Teepublic Shirts
SpreadShirt Apparel

Concepts

Setup

Services

Basic models

_Modelling Steps

Azure Automated Machine Learning

Prerequisites

- Azure subscription
- Azure ML workspace

Data Requirements

- Tabular format
- Value (i.e., label or target column) that you want to predict must be in dataset

ML Algorithms

- **Supervised Learning:** Historical data points with known labels
  - **Regression:** y is numerical
  - **Time Series:** regression with time element, predict feature at a point in time
  - **Classification:** y is categorical

MS AZ AutoML

- No-code ML as a service
- AutoML to build and operate ML solutions in the cloud
- Prepare data, train models, publish predictive services, and monitor their usage.
- Increase data scientist efficiency by automating time consuming processes
- Scale solutions for large volumes of data with pay-as-you-use costs
- Automate time insensitive tasks
- Can add custom code when required (Azure Machine Learning Python SDK)

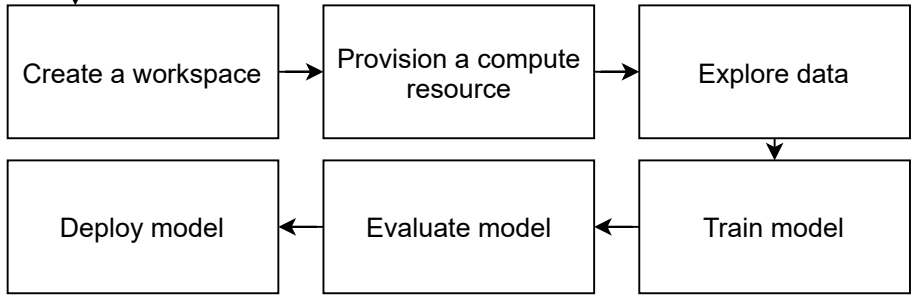

Step 1: Create a workspace

- Via the Azure portal
  - Provide subscription, resource group, workspace name, region, storage account, key vault, application insight and container registry info
- Sign into Azure Machine Learning Studio using Azure directory, subscription & workspace
  - Manage assets and resources via workspace

Step 2: Create a compute resource

- **Compute targets:** cloud-based resources on which to run models and explore data
  - **Compute Instances:** Development workstations that data scientists can use to work with data and models
  - **Computer Clusters:** Scalable clusters of virtual machines for on-demand processing of experiment code
  - **Inference Clusters:** Deployment targets for predictive services on trained models
  - **Attached Compute:** Link to existing resources (VMs and DBs)
- Compute Instances tab
  - Add new compute: provide compute name, VM type and VM size
  - Create compute cluster (while compute instance is pending): provide location, VM priority, VM type, VM size, compute name, min & max nodes, idle seconds, SSH access

Step 3: Explore the Data

- Data is encapsulated in a dataset object - for model training & operations
- Datasets page
  - Create new dataset: provide web url/location, name, dataset type, description, skip validation, file format, delimiter, encoding, column headers, skip rows, multi-line data, schema and confirm details

Step 4: Train the model

- Automated machine learning leverages cloud compute scaling to pre-process and train multiple algorithms in parallel for best performing supervised ML model
  - Classification
  - Regression
  - Time series forecasting
- Automated ML page
  - Create new: provide dataset, name, target column, compute cluster, task type (model), primary (evaluation) metric, explain best model, blocked algorithms, exit criterion (time, score), featurization.
- After completion, Models can be viewed and their scores

Step 5: Evaluate the model

- Review the best model on Details tab & select Algorithm name
- Metric calculated using cross-validation to iteratively tests the trained model with data not present in training data & compare the predicted and actual values
  - Difference between predicted and actual value is the residual and indicates the amount of error in the model (smaller value indicates better model accuracy).
  - Normalize root mean squared error (from residuals)
  - Predicted vs True chart with diagonal trend to show correlation between values
  - Residual Histogram is frequency of residual value ranges, indicates the variance that cannot be explained by the model. Most errors should be clustered around 0.
- Explanations tab > Explanation ID
  - View dashboard of feature importance (Global importance)

Step 6: Deploy model as a service

- As an Azure Container Instance (ACI) or Azure Kubernetes Service (AKS)
  - Production should have AKS with inference cluster compute targets
- On Model tab > Deploy and provide name, description, compute type and enable authentication
- On Endpoint page > select model > consume tab which shows REST Endpoint and Primary Key for your services
  - Notebook page will have model code: file location, file name, file type and overwrite settings.
  - Replace REST endpoint & Primary for service in notebook space provided.
  - Run notebook to complete testing.

ML answer questions that are too complex to answer through manual analysis

- Pieces of code that help people explore, analyze, and find meaning in complex data sets.
- Each algorithm is a finite set of unambiguous step-by-step instructions that a machine can follow to achieve a certain goal.
- In a machine learning model, the goal is to establish or discover patterns that people can use to make predictions or categorize information.
    - Machine learning algorithms use parameters that are based on training data—a subset of data that represents the larger set.
        - As the training data expands to represent the world more realistically, the algorithm calculates more accurate results.

aRtVerse
ART + SCIENCE + NATURE

Compiled by Dr T Oberholster
AI-900 version: October 2021

Website & Blog
aRtVerse Shop
Amazon Books
Teepublic Shirts
SpreadShirt Apparel

ML Algorithms

- **Supervised Learning:** Historical data points with known labels. You have an idea of the final outcome.
    - **Regression:** y is numerical
        - How much/many?
    - **Classification:** y is categorical
        - Is this a/b?
- **Unsupervised Learning:** Unlabelled historical data points. You do not know the outcome.
    - **Clustering:** groups measurements based on shared features
        - What is similar/same/different?

Math & Logic

Models & uses

ML Studio

Modelling Steps

Azure Machine Learning Designer

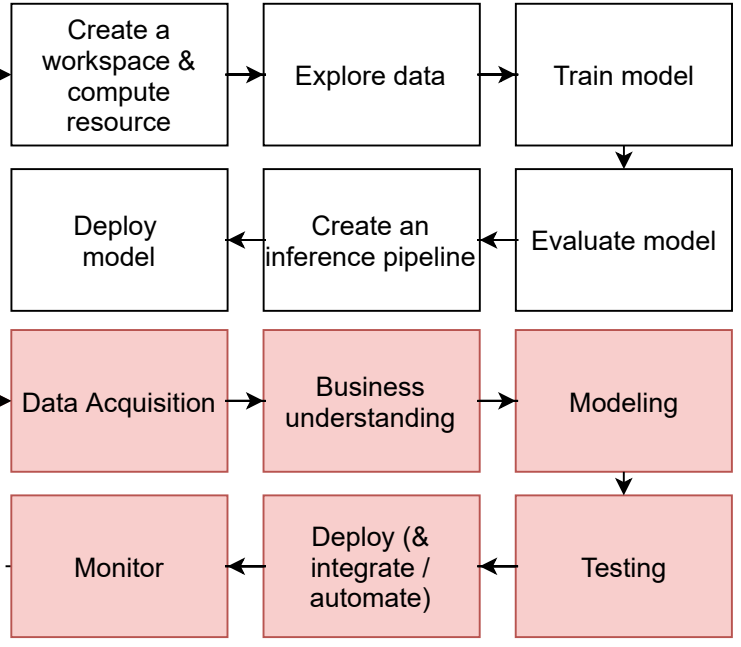

ML Model Lifecycle

Web portal in Azure ML for project authoring

- Notebooks
    - Write and run code in managed Jupyter Notebook servers.
- Azure ML designer
    - Train and deploy machine learning models without writing any code.
    - Drag and drop datasets and components to create ML pipelines.
- Azure AutoML GUI
    - Create automated ML experiments with an easy-to-use interface.
- Data labellling:
    - Efficiently coordinate image labeling or text labeling projects.

Model Building

Components & Assets

- Environment
    - An encapsulation of the environment for training or scoring ML models.
- Experiment
    - A grouping of many runs from a specified script for a workspace.
- Pipeline
    - Create and manage workflows that stitch together machine learning phases.
- Compute resources
    - A machine or set of machines you use to run your training script or host your service deployment
        - Compute instance: VM with multiple ML tools and environments
        - Compute cluster: Cluster of VMs with multi-node scaling capabilities

Components & Assets

- Datasets
    - A reference to the data source location and copy of metadata.
- Datastores
    - Datasets use datastores to securely connect to Azure storage services.
- Models
    - A piece of code that takes an input and produces output
- Endpoints
    - An instantiation of your model into a web service that can be hosted in the cloud.
        - Web service endpoint
        - Real-time endpoint
        - Pipeline endpoint

Step 1: Create a workspace

- Via the Azure portal
  - Provide subscription, resource group, workspace name, region, storage account, key vault, application insight and container registry info
- Sign into Azure Machine Learning Studio using Azure directory, subscription & workspace
  - Manage assets and resources via workspace

Step 2: Create a compute resource

- **Compute targets:** cloud-based resources on which to run models and do data exploration
  - **Compute Instances:** Development workstations that data scientists can use to work with data and models
  - **Computer Clusters:** Scalable clusters of virtual machines for on-demand processing of experiment code
  - **Inference Clusters:** Deployment targets for predictive services on trained models
  - **Attached Compute:** Link to existing resources (VMs and DBs)
- Compute Instances tab
  - Add new compute: provide compute name, VM type and VM size
  - Create compute cluster (while compute instance is pending): provide location, VM priority, VM type, VM size, compute name, min & max nodes, idle seconds, SSH access

Step 3: Explore the Data

- Create a pipeline via Azure ML Studio Designer page
  - Provide: Pipeline name (and date)
  - Select compute target in Settings pane
- Add & explore dataset
  - Drag dataset from **Sample dataset** section (>> to expand panel) onto the canvas (drag-and-drop visual interface)
  - Right-click on dataset, select **Dataset output** (preview data graph icon) on **Output** menu
  - Review schema and distribution histograms
  - Select column header for more details
  - Exclude datasets with high % missingness
  - Close results visualisation window
- Data transformations
  - Expand Data Transformations pane on the left
  - Drag **Select Columns in Dataset** into canvas and connect with Dataset
  - Click into **Select Columns in Dataset**, select **Edit column** & select columns by name using +
  - Add other model and connect them (Clean Missing Data with columns selected, Normalize Data module & transformation method MinMax)
    - Normalization to same scale prevents numeric columns with large values dominating the model outputs
- Run the pipeline
  - Applies data transformations
- View transformed data
  - Select last model & in Settings pane select **Outputs + logs** tab > Visualise
  - Close **results visualisation** window

Regression specific comments

No comments

Classification specific comments

Predicted/Target column has two values (0 or 1), one for each outcome

Clustering specific comments

No comments

Step 4: Train the model

- Create and run training pipeline
  - Train on a subset, rest used to test to compare actual vs predicted labels (model evaluation)
- Data transformation > Split Data & connect to Normalized Data output (left connector)
  - Splitting mode, Fraction of rows, random seed, stratified split
- Model Training > Train Model & connect right input to left output of Split Data
  - Label column

Regression specific comments

- Machine Learning Algorithm > *Regression* > *Linear Regression module* & connect to left input of Train Model
- Testing is done by scoring the validation dataset. Model Scoring & Evaluation > Score Model & connect left input with Train Model output, connect right input with right output from Split Data (carry data and model forward)
- Select Submit to run the pipeline

- Score Model > Outputs & logs > Data outputs > Scored dataset > Preview Data icon
  - Scored labels = predicted values
  - Close results visualisation

Classification specific comments

- Machine Learning Algorithm > *Classification* > *Two-Class Logistic Regression* & connect to left input of Train Model
- Testing is done by scoring the validation dataset. Model Scoring & Evaluation > Score Model & connect left input with Train Model output, connect right input with right output from Split Data (carry data and model forward)
- Select Submit to run the pipeline

- Score Model > Outputs & logs > Data outputs > Scored dataset > Preview Data icon
- Scored labels = predicted values
  - *Scored probabilities between 0 and 1, indicates a positive prediction.*
  - *Probabilities greater than 0.5 indicate 1, and probabilities less than 0.5 indicate 0.*
- Close results visualisation

Clustering specific comments

- Machine Learning Algorithm > *Clustering* > *K-Means Clustering* & connect to left input of Train Model
-
  - *K = number of clusters to create*
    - *Set Number of centroids*
  - *Measurements are treated as multidimensional vectors*
  - *The algorithm initialises K coordinates at randomly selected points (centroids) in n-dimensional space (n number of dimensions in feature vector)*
    - *Feature points are plotted on same dimensional space and each point is assigned to the closest centroid.*
    - *Centroids are moved to the middle of the assigned points (mean distance)*
    - *Reassign points*
    - *Repeat until cluster allocations stabilized or a number of iterations is completed*

- Score Model > Outputs & logs > Data outputs > Scored dataset > Preview Data icon
  - *Assignments column = clusters (0, 1, 2) to which each observation (row) is assigned.*
  - *New columns with distances of assigned points from cluster centers.*
- Close results visualisation

Step 5: Evaluate the model

- Model Scoring & Evaluation > Evaluate Model & connect left input with Score Model
- Select Submit

Regression specific comments

- Compare predicted and actual labels in validation dataset.
- Evaluate Model > Outputs & logs > Data outputs > Evaluation results > Preview Data icon
  - *Mean Absolute Error (MAE): Average difference between actual and predicted values. Lower ~ better model performance*
  - *Root Mean Squared Error (RMSE): Square root of mean squared difference between actual and predicted values. Large values = large variance in residuals*
  - *Relative Squared Error (RSE): Relative squared difference between 0 and 1 of predicted and actual values. Closer to zero ~ better model performance, use for different unit labels*
  - *Relative Absolute Error (RAE): Absolute difference between 0 and 1 of predicted and actual values. Closer to zero ~ better model performance, use for different unit labels*
  - *Coefficient of Determination (R2): R-squared, how much variance between actual and predicted explained by model. Closer to one ~ better model performance*
- Close results visualisation
- More than one regression model can be compared & connected to the evaluation module.

Classification specific comments

- Compare predicted and actual labels in validation dataset.
- Evaluate Model > Outputs & logs > Data outputs > Evaluation results > Preview Data icon
  - *Confusion matrix: Tabulation of predicted and actual value counts for each class. Binary classification models predict 2 values (0 or 1).*

Confusion Matrix		Predicted 1	Predicted 0
Actual	1	True positive	False positive
Actual	0	False negative	True negative

True positive and True negative values should contain the most counts

  - *Accuracy: Ratio of correct predictions (true positive + true negative) to the total number of predictions. How much of the model predictions were correct.*
  - *Precision: Fraction of correctly identified positive cases (True positive / True + False positives). How much of the positive predictions were correct.*
  - *Recall: Fraction of positive cases which are positive (True positives / True positives + False negatives). True positive rate. How much of the actual positives did the model identify.*
  - *F1 Score: Combines Precision and Recall*
  - *Threshold: Slider changes probability scores, moving it to the 0 (Recall = 1) and to 1 (Recall = 0)*

Clustering specific comments

- Evaluate Model > Outputs & logs > Data outputs > Evaluation results > Preview Data icon to view cluster separation metrics
  - *Average Distance to Other Center: Average closeness of each point to all other cluster centers.*
  - *Average Distance to Cluster Center: Average closeness of each point to same cluster center.*
  - *Number of points: assigned to each cluster*
  - *Maximal Distance to Cluster Center: Max distance between each point and cluster centroid. High numbers = high cluster dispersion.*
- *Compare Average Distance to Cluster Center to Maximal Distance to Cluster Center to determine cluster spread*
- Close results visualisation

Regression specific comments

No comments

Classification specific comments

- *<u>ROC curve:</u> Received Operator Characteristic, measures True positive rate (recall) against False positive rate. Large area under the curve = better model performance (AUC). AUC of 0.5 = random change of being correct.*

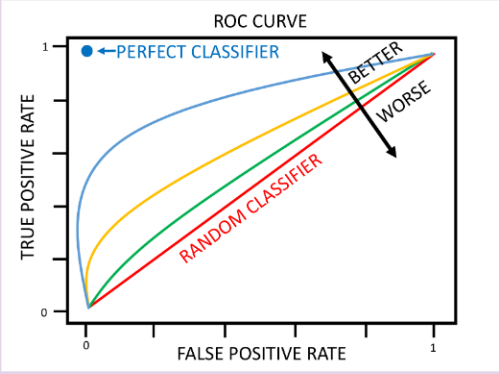

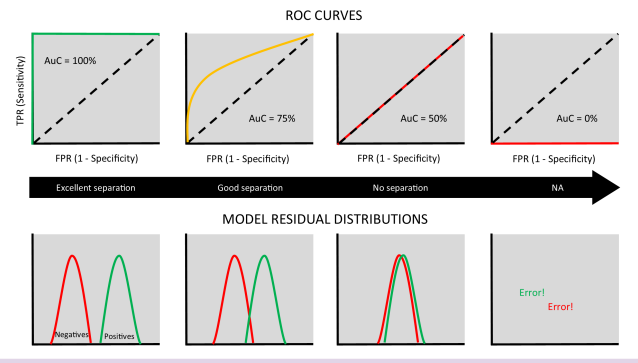

- Close results visualisation
- More than one classification model can be compared & connected to the evaluation module.

Clustering specific comments

No comments

Step 6: Create an inference pipeline

- Perform data transformations on new data & use training model to infer/predict new labels
- Real-time inference pipeline > duplicates training pipeline
  - Rename pipeline
  - Provide Web Service/Enter Data Manually (csv) input for new data & connect to Select Columns in Dataset
  - Remove original raw dataset
  - Remove Evaluation Model module & replace with Execute Python Script (connect left input to Score Model)
    - Replace default values in script with model column names
  - Connect left Execute Python Model output to Web Service Output module (*For Clustering: Assign Data to Cluster)*
- Submit to compute-cluster
  - Execute Python Script > Output & logs > Results dataset (*For Clustering: Assign Data to Clusters > Output & logs > Results dataset)*
  - Close results visualisation

Step 7: Deploy model as a predictive service

- Publish real-time inference pipeline as service for client application
- Deploy
  - Provide name, description & compute type (Azure Contained instance)
  - Test service. Endpoints > select pipeline > Consume
    - REST end point
    - Primary Key
- Create second browser instance of Azure ML Studio > Notebooks (Author) > My Files > New File
  - Provide file location, file name, file type & overwrite settings
  - Check compute instance is running
  - Paste pipeline run code (with one observation's data) into notebook with endpoint & primary key details from previous browser > Consume
- Run notebook & verify return value.

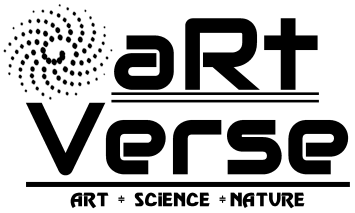

Compiled by Dr T Oberholster
AI-900 version: October 2021

[Website & Blog](#)
[aRtVerse Shop](#)
[Amazon Books](#)
[Teepublic Shirts](#)
[SpreadShirt Apparel](#)

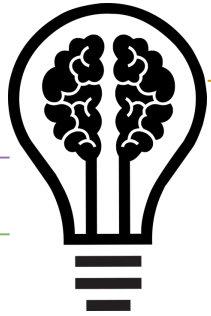

Azure AI Services

AI and ML

Bot Services

MS ML Services

Artificial Intelligence

- Uses complex computer algorithms to process and gain insights from large amounts of data

Machine Learning

- Data science technique that uses existing data to forecast future events and/or outcomes and/or trends

Resposible AI

- Ensures the provisioning of solutions and minimise unintended negative consequences
  - Fairness
  - Reliability & safety
  - Privacy & Security
  - Inclusiveness
  - Transparency
  - Accountability

- Uses Cognitive Services to interact with humans as virtual agents (bots) that understand/reply to questions
- Best for simple, repetitive tasks on websites (book reservations or retrieve information)
- Dialogs through website, email, social media, messaging, phone calls and other channels
- Customer product/services, reservation systems, health care consultants, home automation

- **MS AZ Computer Vision**
  - **QnA Maker Service**
    - Build a knowledge base of question-and-answer pairs, usually with built-in NLP for multiple phrased examples but same semantic meaning.
  - **Azure Bot Service**
    - Develop, publish, and manage bots.
    - Bot Framework to create bot, manage with Bot Service and integrate QnA/LUIS for interactions.

ML as a Service

- Platform for making predictions
- Tools & Services:
  - Define data sources
  - Create pipelines
  - Connect & train models
  - Deploy with API
- Automated model generation & tuning
- Start on local machine & scale out to cloud

ML models

- Develop, train, test & deploy
- Open source tech:
  - Python
  - Jupiter
  - Visual Studio Code
  - Docker
- Automated model generation & tuning
- Start on local machine & scale out to cloud

ML models uses

- Make apps & devices smarter
- Azure Cognitive Services

MS AZ AutoML

- No-code ML as a service
- AutoML to build and operate ML solutions in the cloud
- Prepare data, train models, publish predictive services, and monitor their usage.
- Increase data scientist efficiency by automating time consuming processes
- Scale solutions for large volumes of data with pay-as-you-use costs
- Automate time insensitive tasks
- Can add custom code when required (Azure Machine Learning Python SDK)

MS AZ Machine Learning Designer

- Notebooks
  - Write and run code in managed Jupyter Notebook servers.
- Azure ML designer
  - Train and deploy machine learning models without writing any code.
  - Drag and drop datasets and components to create ML pipelines.
- Azure AutoML GUI
  - Create automated ML experiments with an easy-to-use interface.
- Data labellling:
  - Efficiently coordinate image labeling or text labeling projects.

Compiled by Dr T Oberholster
AI-900 version: October 2021

Website & Blog
aRtVerse Shop
Amazon Books
Teepublic Shirts
SpreadShirt Apparel

Cognitive Services Summary

Azure Computer Vision

- Pre-trained AI models
- **Vision:** Analyse context of Images & videos
- **Speech:** Convert spoken language into text and vice versa
- **Language:** Understand meaning of text
- **Knowledge:** Create resources for apps
- **Search:** Web scale, add free search engine for apps

Language Understanding

- Understands written and spoken language
- Analyse & interpret text with AI
- Interpret spoken language & synthesise speech response
- Translate spoken or written language
- Interpret commands and take actions
- Evaluate different aspects of document/phrase to gain insights into the context of the text.

- **MS AZ Computer Vision**
  - Subset of Analyze Image of Computer Vision
  - Pretrained NLP ML models and services
    - **Text Analytics:** Analyse text documents and extract phrases, detect entities (places, dates, people), and evaluate sentiment (positive/negative)
    - **Translator Text:** Translate +60 languages
    - **Speech:** Recognize, synthesise, and translate speech
    - **Language Understanding Intelligent Service (LUIS):** Train language model to understand spoken or text-based commands

- **MS AZ Computer Vision**
  - Analyse images/videos
  - Extract/Create tags, objects, text & descriptions
  - No ML expertise required

MS AZ Custom Vision
- Train custom image classification and/or object detection models with own images

MS AZ Face
- Build facial detection and recognition solutions

MS AZ Form Recognizer
- Extract information from scanned forms and invoices

Object Detection

- Models predict category/class of object
- Set of inputs (features) to predict labels with probability scores
- Trained on pixel values (feautres) and labels of categorized images

- **MS AZ Computer Vision**
  - Subset of Analyze Image of Computer Vision
  - No need to manage individual servers for ML modelling and classification
  - Auto-scale
  - Azure managed identities
  - Regional resiliency

Facial Detection

- Analyse human faces in images/video
- **Face detection:** ID regions of image with human face
- **Facial analysis:** Return additional information
- **Facial recognition:** Identify known individuals

- **MS AZ Computer Vision**
  - Subset of Analyze Image of Computer Vision
  - **Computer Vision:** Basic analyses (age)
  - **Video Indexer:** Faces in videos
  - **Face:** Widest range of capailities
    - **Identify API:** One-to-many face match
    - **Verify API:** One-to-one face match
    - **Group API:** Similartiy face match

Natural Language Processing

- Intersection of Computer Vision and Natural Language Processing
- Models recognize individual shapes as letters, numbers and punctuation
- Machine Learning Comprehension
  - AI system reads and understands semantic context

- **MS AZ Computer Vision**
  - Subset of Analyze Image of Computer Vision
  - **OCR API:** Text limited
  - **Read API:** Text dominant
  - **Form Recognizer API:** Analyse documents/receipts

Image Classification

- Model recognizes individual objects in an image
- Class of each object in image
- Probability score of classification (confidence of prediction)
- Coordinate bounding box for each object
- Similar to tagging
- Determine multiple instances of same object

- **MS AZ Computer Vision**
  - Subset of Analyze Image of Computer Vision
  - **Detect API**
    - Applies tags in images
  - **Tag API**
    - Include contextual terms (i.e., indoors)